Journal of Information & Communication Technology
Vol. 1, No. 1, (Spring 2007) 11-33

The Effects of Clustering on Performance of the Karachi Stock Exchange

Mansoor-uz-Zafar Dawood[*]
Institute of Business & Technology – BIZTEK, Karachi

Abdul Raouf Khan[*]
College of Computer Science & IT, King Faisal University, Saudi Arabia

Faisal Ijaz[*]
The Karachi Stock Exchange (Guarantee) Limited, Karachi, Pakistan

ABSTRACT

The concept of clustering is not as new as some other concepts in the world of Information Technology. There are many forms of clustering available but the latest revelation in this concept is the database clustering that has transformed the idea of high performance, high availability and seamless scalability. This paper digs deep into the clustering technology concept that the Karachi Stock Exchange has adopted for their mission critical database, which is not only very large in size but also responds to a very large number of transaction each day. The clustering technology implemented at the Karachi Stock Exchange has been developed by Oracle9i Real Application Cluster, which is a feature of Oracle9i Database. This paper identifies performance-related areas of the Oracle clustering technology and collect performance statistics in order, to statistically determine the performance gains for Karachi Stock Exchange from adoption of clustering technology. This research analyzes the effects of performance of the Clustered Database against a Standalone database / single instance Oracle Database and drives the conclusion of how clustering technology is adding to the performance of the biggest pillar in the Pakistani Capital Market i.e. The Karachi Stock Exchange. The tools used for this research were Oracle9i Database, Oracle9i Real Application Clusters (RAC), Oracle10g Internet Application Server and Oracle10g Forms.

Inspec Classification: C6130; C6160; C5220P; C5440

Keywords : Clustering Technology, Performance gains, Oracle 9i database, Stock Exchange

1) INTRODUCTION

The businesses today are becoming more global and require around the clock operation with seamless performance and availability. The Database sizes are becoming huge, storage solution like Storage Area Networks (SAN), Network Attached Storage (NAS) etc. are being used as storage sizes have already peaking the Tera Byte range. Furthermore the

* The material presented by the authors does not necessarily portray the viewpoint of the editors and the management of the Institute of Business and Technology (BIZTEK).

[*]Mansoor-uz-Zafar Dawood : dr.mzdawood@biztek.edu.pk
*Abdul Raouf Khan : raoufkhan@kfu.edu.sa
*Faisal Ijaz : sfaisalalipk@gmail.com

emergence of Data Warehouses and Data Marts technologies has further increased the database size and their criticality.

In a world where storage is becoming cheaper, processing power increasing, server memory increasing Database vendors were left with the challenge to come up with solutions that will compliment these technology / hardware advancements and thus the Clustering technology emerged.

Clustering is a technology in which multiple systems / servers jointly work together but present a single system image to the users. The benefits of clustering are higher performance, large room for scalability, high availability and cost effectiveness.

Karachi Stock Exchange is Pakistan's biggest stock market having the transaction volume in millions each hour. The database record each trade, rate change, the dips and the peaks of the KSE-100 Index and KSE all Share Index.

The Karachi Stock Exchange (KSE) required a similar technology to support its large databases by making them greater in terms of performance and highly available. The Oracle9i Real Application Clusters (RAC) was used for this task. Oracle9i RAC gave the KSE the much need performance and availability feature. The issue remained as to what level has the performance increased. This research is an effort to quantify these performance gains and draw conclusive results from the statistics.

The research starts by identification of areas, which would present data that can be compared to statistics extracted from a single instance database operating in parallel at the KSE along with the Clustered Database. The technology being evaluated is the first implementation of its kind in Pakistan i.e. the Oracle9i Real Application Clusters (RAC). This research involved a tediously long process of collecting the statistics from both the clustered database and evaluating them in the minutest of details.

Karachi Stock Exchange as the biggest stock market of Pakistan encounters about 2.3 Million direct and in direct transaction on its database.

The availability of the system was another challenge for the Karachi Stock Exchange. Keeping the systems running even if a failure occurs and recover to a state of normality in the shortest possible time span was a priority of KSE.

To cater these issues, KSE referred to the Clustered Database technology for the implementation, following were the problems at hand:

- Evaluate the performance of a Clustered Database against a standalone / single instance database
- Finding out what parameter might be necessary in determining the performance
- Making the system that will provide a mechanism for the collection of Statistics
- Gathering Statistics at certain intervals
- Evaluating these statistics to draw a conclusion

2) KARACHI STOCK EXCHANGE RESEARCH SPECIFICATIONS

2.1) PERFORMANCE PARAMETERS

The Karachi Stock Exchange has a very complex computing environment running multiple applications made on various different platforms e.g. Java, Oracle Developer and Visual

C++. All the applications play a part in trading and settlement activities that take place at the exchange and thus undergo an intense transaction volume.

To gauge performance of the Clustered Database, identification of parameters was required. These parameters should be found from each aspect of the working of a clustered system e.g. Memory usage, Processing, space allocation, objecting events & locking, and user session parameter.

An intense research was carried out in two different aspects of Oracle Server computing, firstly the internal architecture of Oracle9i Database Management System and secondly the Oracle Database or the physical storage. The analysis phase of the research highlighted the two sources from which the performance statistics has been extracted. These sources are:

- Two Clustered Database Instances
- Stand Alone Instance

2.2) PARAMETERS CHOSEN FOR ANALYSIS

The analysis of the KSE study revealed about 100 parameters which provided enough data for comparison of performance in an environment with a cluster database and an environment without a clustered database. Amongst these 100 parameters, 18 parameters were selected on the basis of:

- Continuous short-listing of parameters after study of Oracle documentation
- Impact of the parameters in terms of research concentration.
- Importance of parameters in terms of performance based on professional advice from RDBMS professionals

The selected eighteen parameters are:

Buffer No wait Percentage, Buffer Hit Percentage, Library Hit Percentage, Execute to Parse Percentage, Parse CPU to Parse Elapsed Percentage, Redo No Wait Percentage, In-Memory Sort Percentage, Soft Parse Percentage, Latch Hit Percentage, Percentage Non-Parse CPU, CPU Time, Control File Sequential Read, DB File Sequential Read, Log File Sync, Log File Parallel Write, Log File Parallel Write, PGA Cache Hit Percentage, and Parse Count (Hard)

2.3) COLLECTION OF DATA

After identification of performance parameters or Key Performance Indicators (KPI), the next most important phase of the research was collection of statistics. The parameter values were stored in Oracle database's data dictionary tables. Some of these tables hold data periodically and some had data which was updated from time to time. All these tables were dynamic performance tables i.e. the data in them reflected the current situation of the database operations.

A detailed script was written that extracted data from these tables at regular intervals. This script was created using PL/SQL (Procedural Language / Structured Query Language) and SQL (Structured Query Language).

The script was scheduled to collect data at a specified time with a specified interval. The data collected through this script was stored in a text file with data appended to it. The parameter was scheduled using Oracle database's native package for job scheduling called DBMS_JOB.

2.4) DATA COLLECTION INTERVAL

The Karachi Stock Exchange is a mission critical setup with huge volume of transaction involving large sums of monetary value. Hence the system undergoes a great deal of stress. The scripts written gathering performance statistics were also very resource intensive. Therefore in consultation with the KSE IT management following point with regards to statistics gathering were agreed:

Table 1
Data Collection Characteristics

Statistics gathering interval	2 Weeks			
Number of readings	8			
Dates for statistics gather	21-Feb-05 18-Apr-05	7-Mar-05 2-May-05	21-Mar-05 16-May-05	4-Apr-05 30-May-05
Time for execution of script	11:15 AM			
Total duration of statistics	4 Months			

3) ANALYSIS OF PERFORMANCE PARAMETERS

Following is the detailed analysis of each of the selected parameters:

1. Buffer No Wait Percentage

Parameter Name	**Buffer No wait percentage**							
Parameter Description	The Buffer No wait percentage refers to the data blocks that come from hard disk and looking for free space in database buffer cache. The parameter refers to the percentage of blocks that found space in the memory without any waits.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	98.76	99.1	96.83	99.7	100	99.57	100	98.76
Instance 2	94.67	97.13	99.99	99.7	88.15	98.17	100	91.42
Stand Alone	76.91	80.4	77.32	69.49	83.8	87.1	71.05	70.9

Figure 1
Buffer No wait percentage Graph

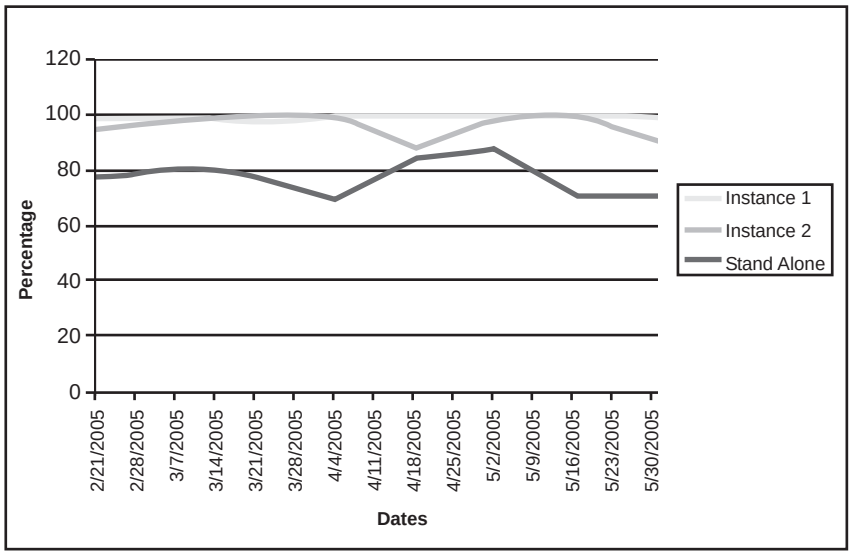

Result Analysis

The Buffer No wait percentage parameter clearly states that both Instance 1 and Instance 2 reading are much more close to the target 100% then the standalone instance, With the standalone instance the performance is very variable. The statistics clearly reveal the Buffer No wait percentages are much higher for RAC Instance as compared to Standalone instance.

2. Buffer Hit Percentage

Parameter Name	**Buffer Hit percentage**							
Parameter Description	The Buffer Hit percentage parameter refers to the percentage of Data Blocks required by a query that were found in a memory and did not require a read from the Hard Disk.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	99.8	99.23	95.58	98.3	99.22	97.81	99.11	99.8
Instance 2	96.41	100	99.76	99.8	96.83	92.63	99	94.12
Stand Alone	82.75	86.13	79.95	78.66	74.52	85.11	90.2	71.2

Figure 2
Buffer Hit percentage Graph

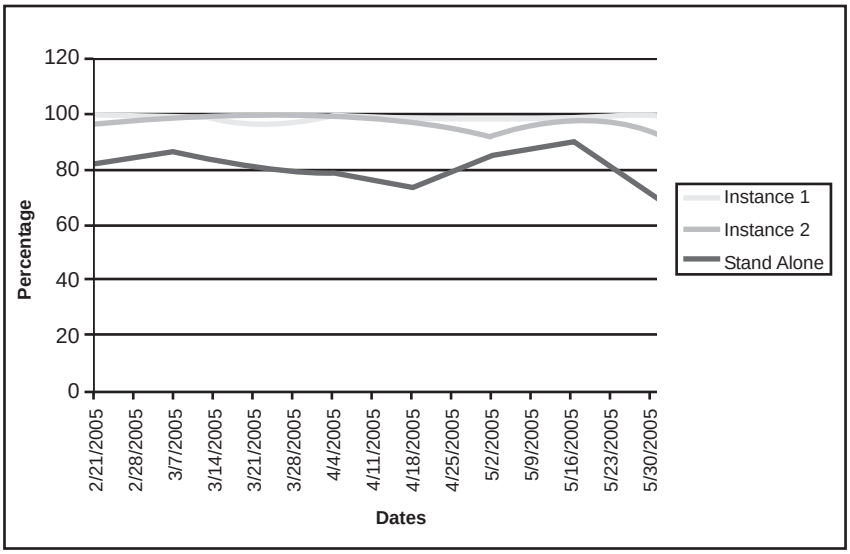

Result Analysis

The Buffer hit percentage is another parameter where both the clustered instances are showing values very close to 100% as compared to its standalone counterpart. The Buffer Hits percentage values for Stand Alone System are much lower simply because of the larger size of the DB Buffer and features in the Cluster database like Cache Fusion which enables sharing in caches of Clustered Databases.

3. Library Hit percentage

Parameter Name	**Library Hit percentage**							
Parameter Description	The Library Hit percentage parameter refers to the percentage of queries that do find there parse plans in the memory rather then being parsed again.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	100	96.78	93.337	100	89.32	100	100	100
Instance 2	99.26	100	100	100	99.74	99.29	97.89	99.03
Stand Alone	82.72	89.21	80.74	79.83	90.2	87.2	76.3	77.52

Figure 3
Library Hit percentage Graph

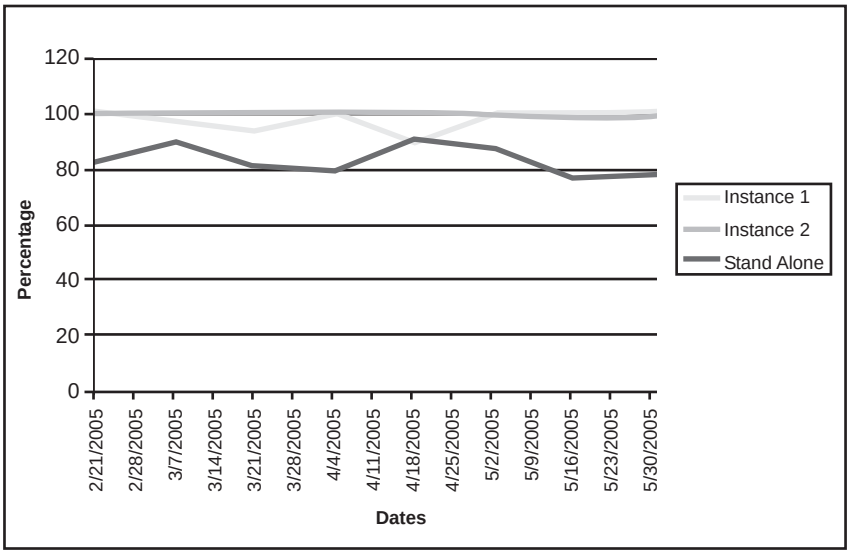

Result Analysis

Library Cache is a place in the memory where parse plans of queries are kept. The Library cache of a RAC database becomes much larger in size as it enables database to share Library cache yet projecting a single system image. The Library Hit percentage in both the RAC database instance is very much better to the target value of 100% as compared to the stand alone instance where the Hit percentage is much less.

4. Execute to parse percentage

Parameter Name	**Execute to Parse percentage**							
Parameter Description	The Execute to Parse percentage parameters refers to the percentage of ratio between the total number of queries executed and the number of queries required a parse i.e. those parse plans did not exist in the Library Cache.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	98.1	99.37	98.19	91.02	92.39	91.1	96.42	88.19
Instance 2	100	94.33	81.26	93.32	91.31	91.53	93.13	89.24
Stand Alone	50.3	69.4	72.1	70.3	71	65.08	69.85	58.72

Figure 4
Execute To Parse percentage Graph

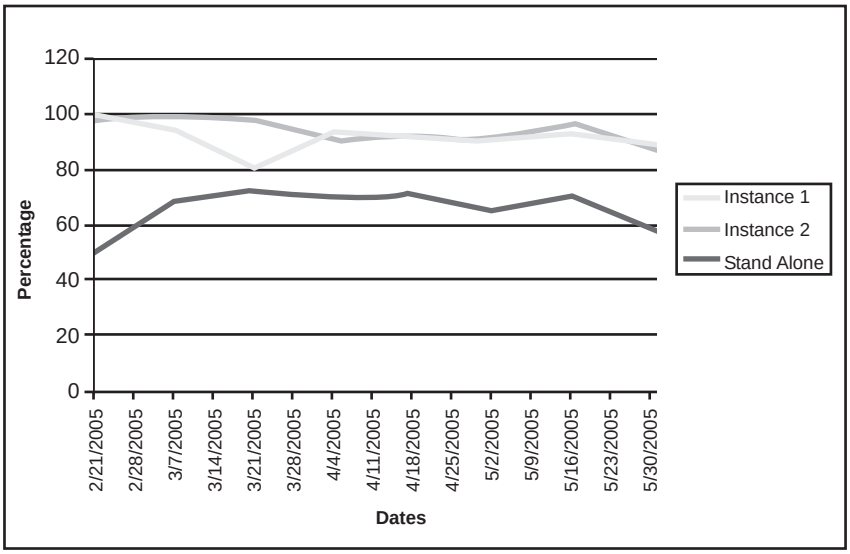

Result Analysis

The execute to parse percentage are consistently high in the clustered database as the area which holds these parse plans is larger as it is utilizing Random Access Memory from two different machines and sharing information in each through Cache Fusion Technology. Thus the Performance of RAC Systems in the aspect of computing is much superior to Stand alone systems.

5. Parse CPU to Parse Elapsed percentage

Parameter Name	**Parse CPU to Parse Elapsed percentage**							
Parameter Description	This parameter refers to queries that the time CPU spent in parsing queries as compared to total time spent in parsing.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	99.33	99.67	97.63	100	98.73	88.65	92.83	91.44
Instance 2	90.73	86.97	96.91	93.67	93.67	96.51	93.06	93.7
Stand Alone	71.96	77.92	68.3	55.02	54.59	70.03	62.39	60.2

Figure 5
Parse CPU to Parse Elapsed percentage Graph

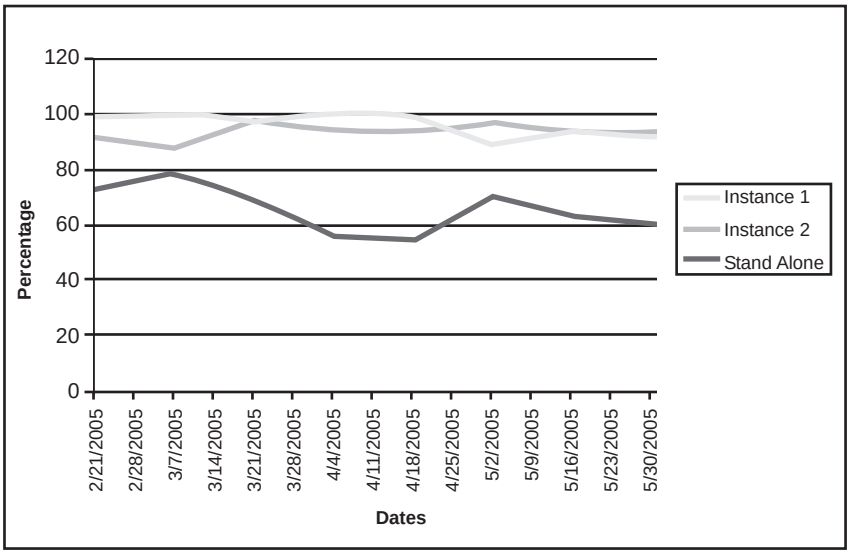

Result Analysis

Both the Clustered Instances have Parse CPU times in high 90's which is a very good level for enterprise environments. The Stand Alone system is again the Parse CPU values are much lower signifying that the parse time for a general query are exceeding a threshold value set by the Oracle optimizer as the CPU is over burdened.

6. Redo No Wait percentage

Parameter Name	**Redo No Wait percentage**							
Parameter Description	The Redo No Wait percentage parameter refers to Redo entries /blocks that did not need any waiting when space was required for data to be written in the memory (Redo Log Buffer).							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	99.1	99.1	99.1	99.1	99.1	99.1	99.1	99.1
Instance 2	100	100	100	98.97	98.41	97.38	90.07	94.41
Stand Alone	83.52	80	79.99	84.2	83.77	78.03	79.3	80.91

Figure 6
Redo No Wait percentage Graph

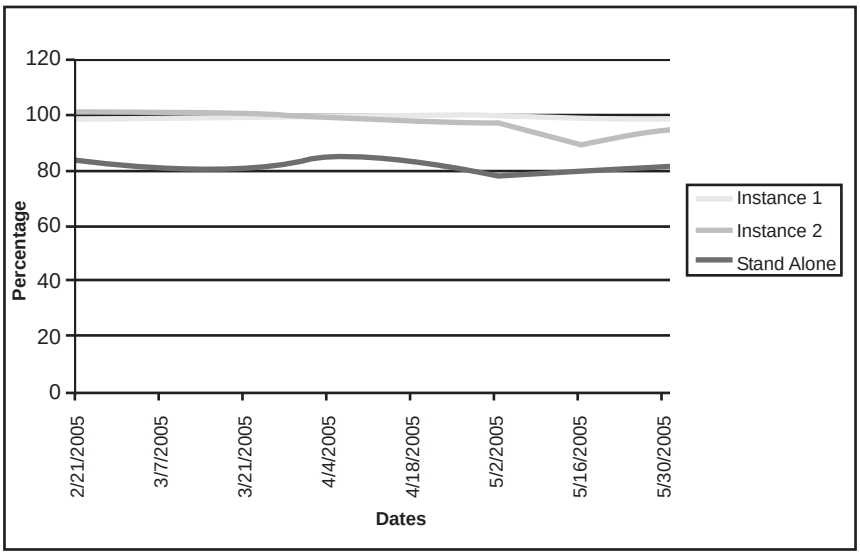

Result Analysis

The Redo No wait values are again lower of the Stand Alone System as its redo log buffer is smaller as compared to Cluster database. It is also a fact that because load is being balanced between the nodes therefore the Redo is also shared in RAC system and hence less waits occur.

7. In-memory Sort percentage

Parameter Name	**In-memory Sort percentage**							
Parameter Description	This parameter refers to the sorting space that is required for queries that need sorting area. Each Oracle session has a small space in memory. If that space is adequate sorting is done in the memory which extremely as compared to sorts that need space on the hard disk for sorting.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	100	98.75	100	90.1	97.13	99.61	98.67	100
Instance 2	100	98.7	94.261	100	94.2	97.32	100	100
Stand Alone	74.23	89.02	76.49	77.13	81.1	74.79	79.2	70.06

Figure 7
In-memory Sort percentage Graph

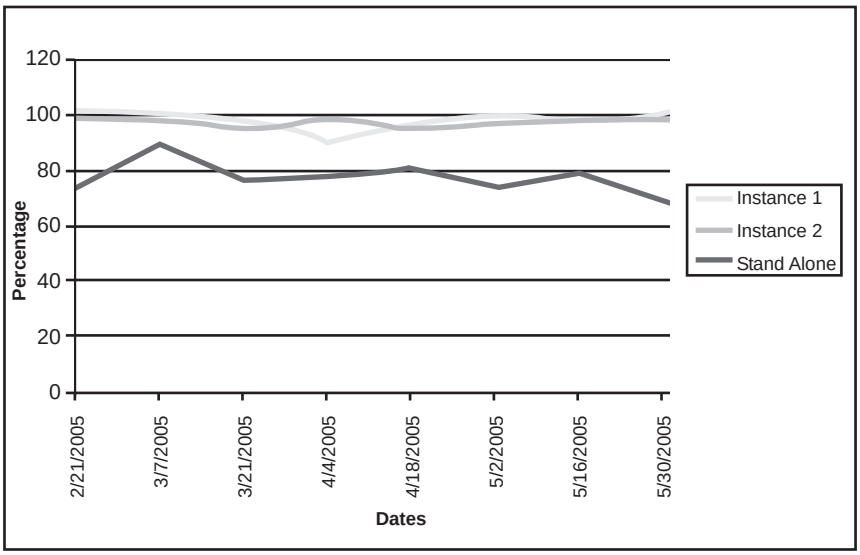

Result Analysis

In-memory sorts are preferred as compared to disk sorts. The Clustered Database has larger memory and as connections are shared between nodes PGA are also split between nodes and hence the performance of Clustered Database is much higher as compared to standalone Database.

8. Soft Parse percentage

Parameter Name	**Soft Parse percentage**							
Parameter Description	Soft Parse percentage refers to the queries that did not require parsing i.e. their parse plans already existed in the memory. The parameter is the percentage of total number of parses against software parses.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	99.6	99.62	100	99.6	98.76	95.62	97.32	95.16
Instance 2	99.2	99.6	99.98	98.11	98.76	95.66	97.31	99.6
Stand Alone	86.3	78.24	57.33	67.39	63.25	60.1	64.05	66.42

Figure 8
Soft Parse percentage Graph

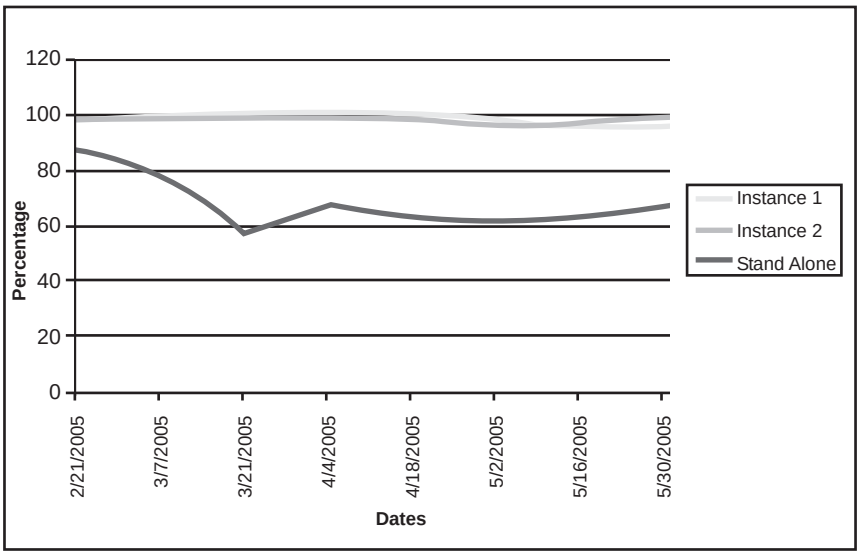

Result Analysis

The Soft Parse percentage of Clustered Databases is very close to the 100% target because most of the queries are finding there parse plans in the memory as compared to single instance / Stand Alone Database where soft parsing is very low. Hence Clustered database is performing much better.

9. Latch Hit percentage

Parameter Name	**Latch Hit percentage**							
Parameter Description	The Latch Hit percentage metric is the ratio of the total number of latch misses to the number of latch gets for all latches. A low value for this parameter indicates a latching problem, whereas a high value is generally good.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	93	90.18	98.01	96.12	100	99.18	99.45	94.3
Instance 2	98.11	99.42	96.28	98.44	90.81	100	99.52	98.74
Stand Alone	76.3	80.12	83.21	74	84.17	73.2	77.14	76.93

Figure 9
Latch Hit percentage Graph

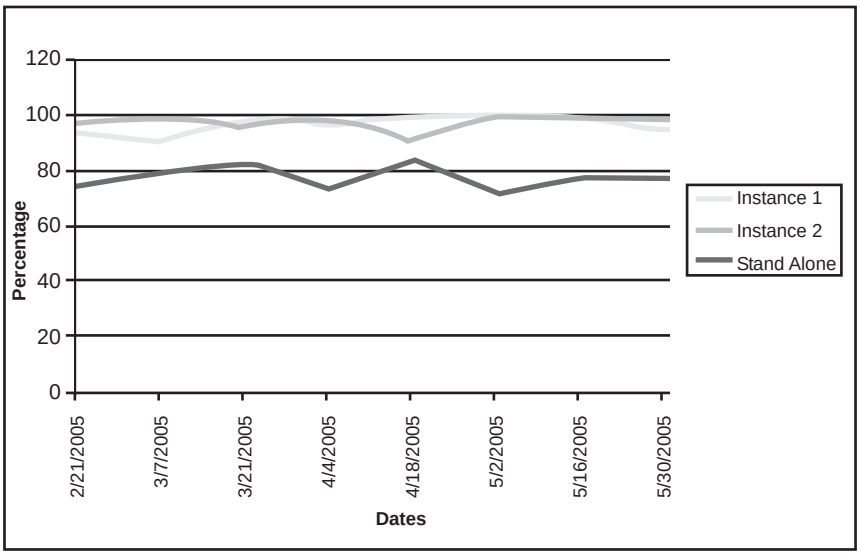

Result Analysis

The Latch hit percentage in Clustered systems is once again very high thus signifying better performance as compared to its single instance counterpart.

10. Percentage Non-Parse CPU

Parameter Name	**Percentage Non-Parse CPU**							
Parameter Description	The percentage Non-Parse CPU refers to the percentage of CPU used by queries that did not require any parsing against the CPU consumed by queries that required parsing.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	79.68	86.54	81.11	82.66	86.54	88.97	94.6	89.32
Instance 2	99.54	93.12	99.03	82.66	86.54	79.68	89.97	86.54
Stand Alone	60.54	63.21	67.95	76.43	70.36	69.64	71.2	74.85

Figure 10
Percentage Non-Parse CPU Graph

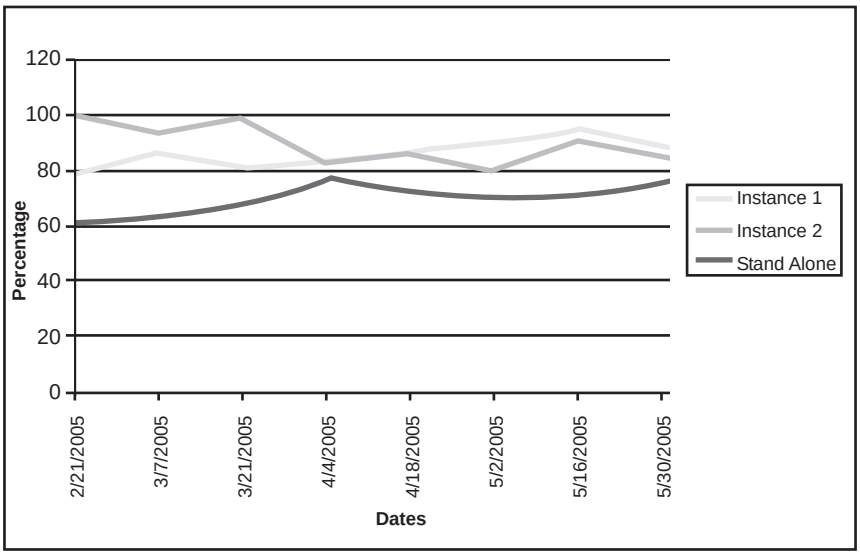

Result Analysis

The percentage Non-Parse CPU for Clustered Databases is better performing as cluster database's have more CPU / processing power available as compared to non-clustered instance.

11. CPU Time

Parameter Name	**CPU Time**							
Parameter Description	This parameter refers to total time that was required by the CPU to parse queries.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	0.11	0.19	0.07	0	0.87	0.02	0	0.65
Instance 2	0.23	0.19	0.38	0	0.65	0.12	0	0.65
Stand Alone	1.88	2.05	2.1	2.41	1.73	2.05	1.95	2.52

Figure 11
CPU Time Graph

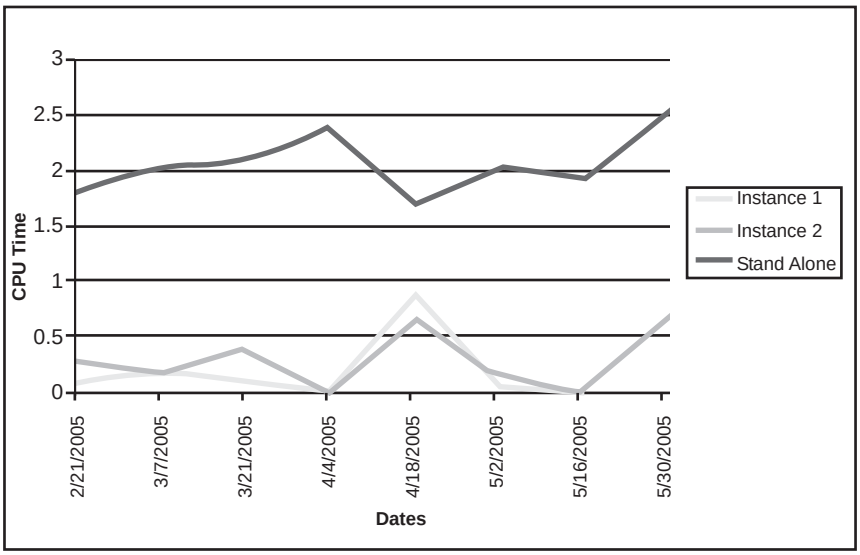

Result Analysis

If Parse plans are phased out of Shared Pool, queries have to be reparsed which takes longer CPU time. Cluster databases have consumed far less CPU time as compared to its Stand alone counterpart as its Memory which holds parse plans is larger and CPU power is greater.

12. Control File Sequential Read

Parameter Name	**Control File Sequential Read**							
Parameter Description	The number of time a control file was required for sequential reading as compared to the number of time it was not available for reads.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	0.59	0	0.51	0	3.34	0.91	4.34	7.34
Instance 2	0.02	0	0.59	0.01	0.34	0.53	0.05	0.34
Stand Alone	1.72	1.66	1.71	2.24	2.091	2.05	2.33	2

Figure 12
Control File Sequential Read Graph

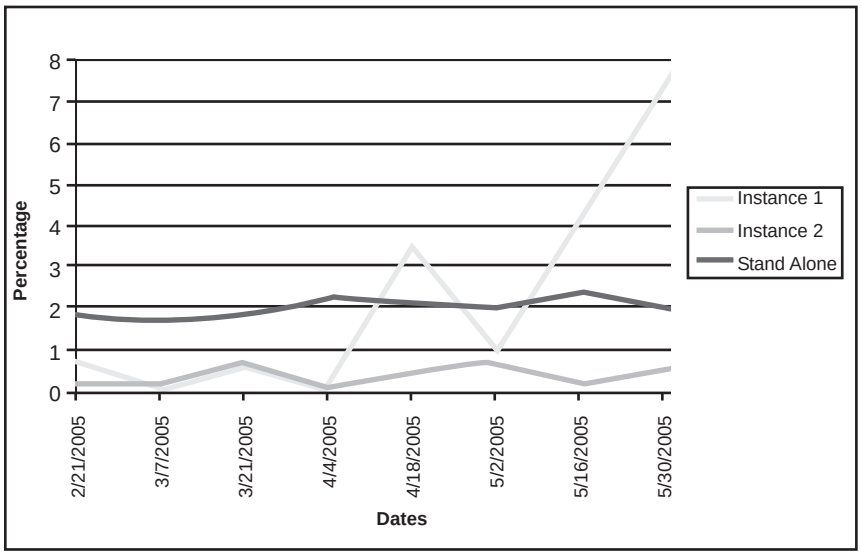

Result Analysis

Control file is a file which controls the database. Its read time should be minimal as it costs heavily to read this file. Control file of Instance 2 is far less for instance 1 which at times as even crossed even its stand alone counter part. A reason for this could be some high resource activity at the instance 1.

13. DB File Sequential Read

Parameter Name	**DB File Sequential Read**							
Parameter Description	The number of time a Database file was required for sequential reading as compared to the number of time it was not available for reads.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	1.69	0.04	1.02	0	2.29	0.04	2.04	4.04
Instance 2	0.53	0.28	0.56	0.36	0.29	0.24	0	4.54
Stand Alone	1.58	1.43	1.5	1.61	1.33	1.042	1.87	3.15

Figure 13
DB File Sequential Read Graph

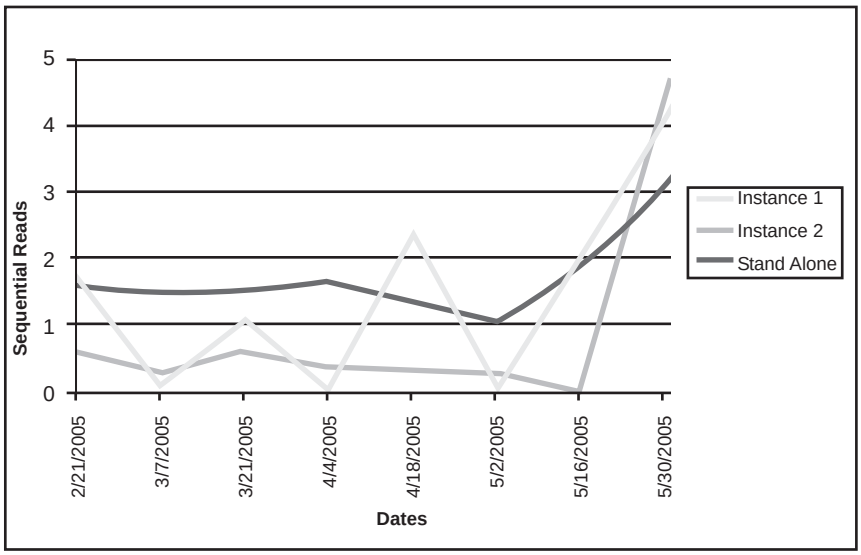

Result Analysis

DB file (Data File) when read sequentially degrade performance. The Clustered instance have shown a trend which is similar to the Standalone instance at times but generally the cluster instances are consuming much less time then standalone instance thus performing better.

14. Log File Sync

Parameter Name	**Log File Sync**							
Parameter Description	The number of times log files were available for sequential writing as compared to the number of time it was not available for writes.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	0.21	0.06	0.911	0	0.04	0	3.06	0.06
Instance 2	0	0.06	0.99	0	0.032	0	0.58	0.26
Stand Alone	2.3	2.51	2.14	2.57	1.82	1.52	1.04	1.32

Figure 14
Log File Sync

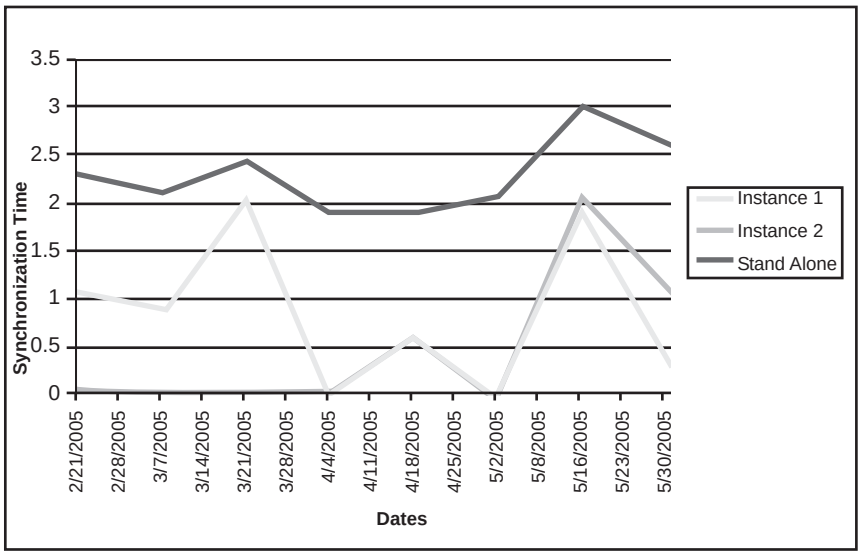

Result Analysis

Log file synchronization time is also a Key performance indicator as far as database technology goes. The Cluster database instances have performed much better in this instance as there writing capability in the memory is greater then a stand alone instance.

15. Log File Parallel Write

Parameter Name	**Log File Parallel Write**							
Parameter Description	The number of times log files were available for sequential writing as compared to the number of time it was not available for writes.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	0.21	0.06	0.911	0	0.04	0	3.06	0.06
Instance 2	0	0.06	0.99	0	0.032	0	0.58	0.26
Stand Alone	2.3	2.51	2.14	2.57	1.82	1.52	1.04	1.32

Figure 15
Log File Parallel Write

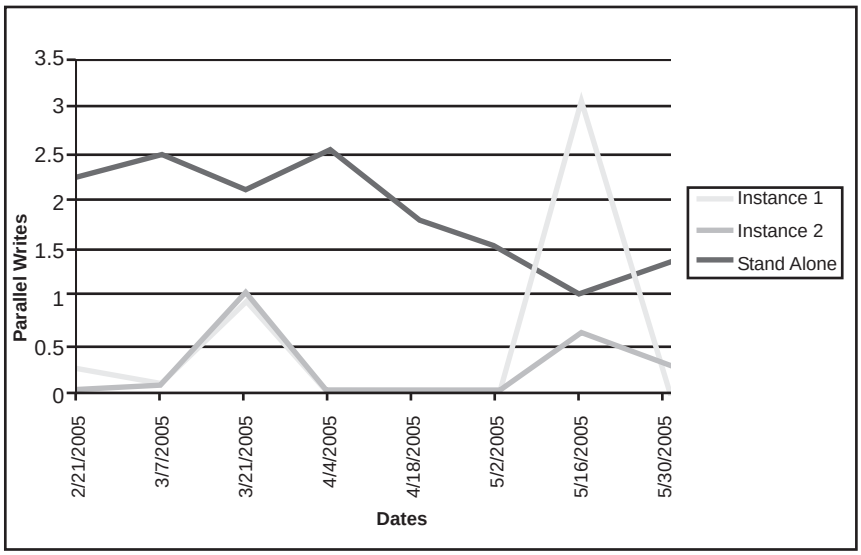

Result Analysis

Parallel writing is very expensive and consumes a lot of CPU time. The Log file parallel write parameter is showing that the clustered database instance through there enhanced writing capability are showing higher performance as compared to non-clustered instance as its parallel writing capability is faster.

16. CPU Used By This Session

Parameter Name	**CPU Used By This Session**							
Parameter Description	The total amount of CPU time that was used by this session that was gathering statistics for performance comparison.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	162	115	199	107	156	128	193	152
Instance 2	102	152	132	127	133	169	201	128
Stand Alone	487	431	402	514	497	522	412	385

Figure 16
CPU Used By This Session

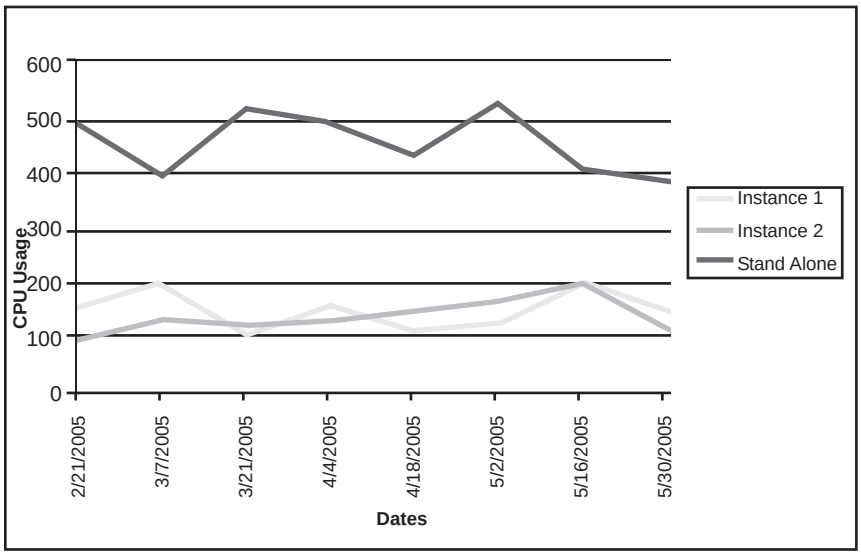

Result Analysis

This parameter is a general indicator of the CPU which this session collecting data for performance measurement used. As with RAC instance the CPU usage is split between nodes therefore the usage of CPU is higher in a non-clustered database then a clustered database.

17. PGA Cache Hit percentage

Parameter Name	**PGA Cache Hit percentage**							
Parameter Description	The percentage of hits that were recorded in the PGA. The ratio is between hits and misses from the Cache.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	100	100	99.42	99.36	100	100	98.94	99.48
Instance 2	100	100	100	99.89	100	95.6	100	98.72
Stand Alone	68.5	65.41	70.01	71.33	73.14	72.51	78.51	76.35

Figure 17
PGA Cache Hit percentage

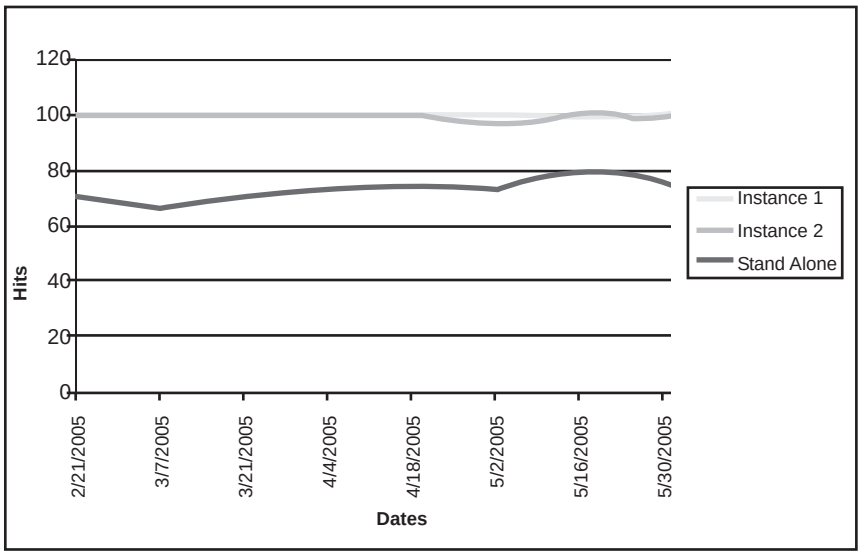

Result Analysis

The PGA is a memory component which is assigned to each process. The Cache hits in a clustered instance are much higher as the memory allocated to the cluster instances is greater because of greater number of servers. Hence the cluster databases have larger number of hits.

18. Parse Count (Hard)

Parameter Name	**Parse Count (Hard)**							
Parameter Description	This parameter refers to queries that found enough space in the memory to be parsed.							
Parameter Values								
	21-Feb	7-Mar	21-Mar	4-Apr	18-Apr	2-May	16-May	30-May
Instance 1	16	12	2	1	8	17	19	12
Instance 2	2	8	5	11	4	18	6	4
Stand Alone	31	42	38	42	34	33	47	31

Figure 18
Parse Count (Hard) %

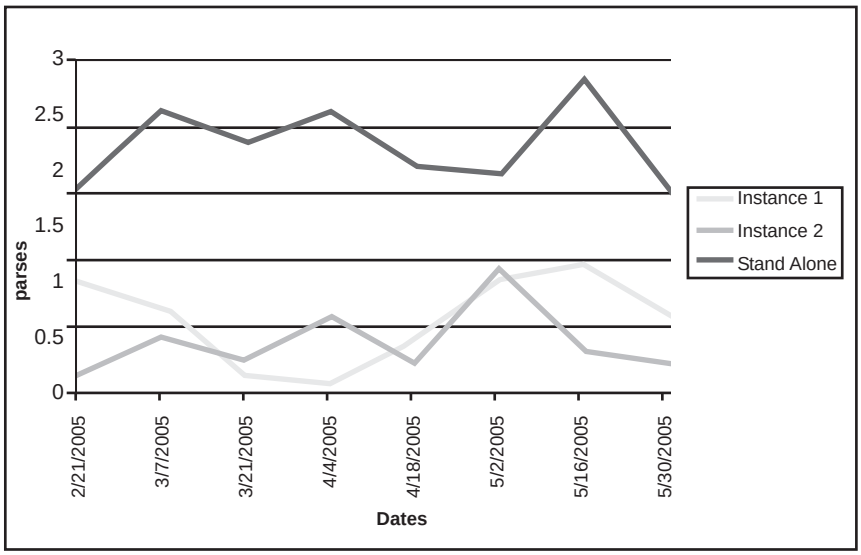

Result Analysis

The clustered database instances have less physical writes due to larger and more efficient usage of memory which requires less writes from the hard disk. Thus the RAC systems performance level is higher as compared to non-clustered instance.

3) CONCLUSION

Clustering Technology is revolutionizing the world of enterprise computing where the volume of data and the criticality of business processes is increasing day in and day out. Clustering technology has been proving itself in environments where high performance is not only required but is critical to daily business operations, and its live example can be seen at the Karachi Stock Exchange.

After detailed analysis of parameters, it is clear that the clustered database instances have performed better and clustering technology creates transparency among the multiple systems.

This conclusion is drawn from data, that was collected over a period of four months through scripts / code written and executed in parallel against both the cluster database instances and non clustered database instances.

The performance gains to Karachi Stock Exchange (KSE) in all aspects of computing are higher with the reduction in cost, as clustered databases can be enhanced with security to a large level in a high availability environment allowing smooth and efficient running of stock market operations.

REFERENCES

AULT, M., TUMMA M., BURLESDON, D, (2003) “Oracle Real Application Clusters Configuration and Internals” Rampant TechPress

VALLATH M (2003) “Oracle Real Application Clusters”, Digital Press
POWELL GAVIN GT (2003) “Oracle High Performance Tuning for 9i and 10g” Digital Press
http://www.oracle.com/technology Managing Oracle Real Application Clusters, Oracle Corporation., , (Last accessed on March 11, 2005)
http://www.oracle.com/technology/products/database/clustering Reality Behind Real Application Clusters Marketing Messages, Oracle Corporation., (Last accessed April 19, 2005)
http://www.oracle.com/technology/deploy/perf/pdf/FederatedvsClustered.pdf Database Architecture: Federated Vs Clustered, Oracle Corporation. (Last accessed July 11, 2005)